

## Conclusions

- CALS achieves a higher arithmetic intensity than a sequence of CP-ALS invocations without numerically affecting the computation and therefore completes a set of decompositions faster.




Fig. 2: Times for CALS, ALS and Tensor Toolbox on a real dataset.

- CALS makes offloading to a GPU worthwhile by increasing the granularity of the central MTTKRP operation, which further increases the speed.




Fig. 3: Times for CALS (12 & 24 threads, CUDA) on artificial data.